DOCUMENT RESUME

ED 059 239 TM 001 049

AUTHOR	Novick, Melvin R.
TITLE	Bayesian Computer-Assisted Data Analysis.
INSTITUTION	American Coll. Testing Program, Iowa City, Iowa. Research and Development Div.
REPORT NO	ACT-TB-3
PUB DATE	[71]
NOTE	27p.
EDRS PRICE	MF-$0.65 HC-$3.29
DESCRIPTORS	*Bayesian Statistics; *Computer Oriented Programs; *Data Analysis; *Interaction; Models; *Nonparametric Statistics; Norms; Statistical Analysis

ABSTRACT

An interactive computer-based system for assisting investigators in the use of Bayesian analysis using the two parameter normal model is described. An important feature of this program is that it interacts with the investigator in the English language; he need not be familiar with computer languages or with the internal workings of the computer. (MS)

ED 059239

ACT TECHNICAL BULLETIN NO. 3

BAYESIAN COMPUTER-ASSISTED DATA ANALYSIS

by

Melvin R. Novick

U.S. DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
OFFICE OF EDUCATION
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL OFFICE OF EDUCATION POSITION OR POLICY.

TM 001 049

The Research and Development Division

The American College Testing Program

P. O. Box 168, Iowa City, Iowa 52240

BAYESIAN
COMPUTER-ASSISTED DATA ANALYSIS

by

Melvin R. Novick

Introduction

Few scientific investigators can maintain expertise in statistical methods. They may have had substantial training in, understanding of, and competence in statistical methods; but they are unlikely to exercise these skills often enough to maintain them at a high level of proficiency. These investigators can use and are typically receptive to guidance in their statistical work provided they remain in control of their own analyses. Investigators with lesser statistical skills will benefit from even more directive guidance through the maze of detail required of good statistical practice. For all investigators, the tedium of computation or alternatively, the maintenance of esoteric computer expertise, is a regrettable hindrance to their function of extracting meaning from data.

The system described here is an interactive computer-based system for assisting investigators on a step-by-step basis in the use of a particular analytic tool--Bayesian analysis using the two parameter normal model. The example is meant to be suggestive of the kinds of computer-assisted data analysis programs that can be developed for use by scientific investigators. Programs such as these can also be used in the classroom and laboratory for teaching purposes, but beyond this, they can be used by the practicing scientist in his day-to-day work.

The present program begins by assisting the investigator in the specification of his prior distribution of the unknown standard deviation and the unknown mean as a particular member of an appropriate extended natural conjugate class. As defined here, this class is somewhat broader than usually defined and, hence, more likely to accurately reflect prior knowledge or beliefs. For each parameter, the investigator is asked to specify three pieces of information about his prior information while, in fact, only two of these are necessary to determine this distribution. The requirement for three pieces of information typically evokes an inconsistency on the part of the investigator. The resolution of this inconsistency through sequential adjustment of judgments results in the statement of a coherent prior distribution. Throughout the interrogation, the investigator remains in control of his specification of prior information and is given repeated opportunity to revise his judgments, explore further, and reject simplifying statements.

An important feature of this program is that it interacts with the investigator in the English language. The investigator need not be familiar with computer languages or with the internal workings of the computer. He need only learn how to sign-in and sign-off the terminal and to make simple alphabetic and numeric responses.

Some investigators, however, will have a somewhat more fully developed understanding of the workings of the interactive program. For these investigators, the interactive program provides some techniques for speeding up the analysis. An investigator who has used a program a number of times will probably feel that he needs a decreasing amount of guidance. Built-in "interrupt" and "go to" functions make it possible for him to proceed more quickly as his skills increase. For example, the investigator who is sure of his prior distribution can go directly to the posterior analysis.

The program also makes it possible for the investigator to consider easily a host of subsidiary questions that might be raised. He can, for example, get some idea of whether his contemplated sample size is adequate by pretending that he has obtained certain results and then seeing whether or not the posterior distribution given these fictitious data is tight enough for his needs. He can also similarly determine what sample size would be required to sufficiently modify the beliefs of a person holding different prior beliefs than his. A typical result of this kind of analysis is the realization that rather large samples are usually required to modify some prior beliefs--samples that may be larger than are economically feasible to obtain.

We would emphasize that the program described here, while useful in its present form, should be thought of primarily as the prototype of more detailed, sophisticated, and complete programs. Our purpose in distributing this preliminary report is to encourage other investigators to become interested in some of the problems that evolve from this approach to interactive computing.

The Problem

We shall illustrate the use of the CADA program for a specific problem that is typical of a type of problem that arises with high frequency. An investigator needs to get a pretty good idea of the level of attainment of college-bound graduating high school students in the state of Iowa. He needs this information quickly. To provide some information, he looks at some data obtained from a published report with the hope that he will not have to get further data.

The information he has consists of scores on The American College Test for twenty students who participated in a statewide survey. He has examined the process by which these students were selected and finds no evidence to challenge the assertion that this is a random sample of current college-bound students. More specifically, he has no information to distinguish these students from other students in the population.

The background information available to him is as follows: He knows that nationally the mean ACT score is 20, and the standard deviation is 5. He knows that the state of Iowa has a very high literacy rate and a very high per pupil expenditure on schools. This leads him to expect a mean value greater than 20. Also, as a statistician, he expects, on the average, a greater homogeneity in one state than across states.

Undoubtedly, some investigators will have more specific prior information, and thus, we must expect that different investigators will wish to use different prior distributions. Our task is to assist each investigator in quantifying his prior beliefs. We shall follow our investigator through the CADA program, annotating the ouput where some explanation seems useful. We shall then follow a second and, substantively better informed investigator through the same program. Finally, we investigate the size of sample that might be required to get consensus from the two investigators.

BAYESIAN DATA ANALYSIS WITH THE NORMAL MODEL WITH UNKNOWN MEAN AND STANDARD DEVIATION

This program will assist you in specifying a prior distribution on the standard deviation and a prior conditional distribution for the mean given the standard deviation. After arriving at a satisfactory specification of priors, sample data from your experiment will be used to compute the relevant posterior distributions. To do the posterior analysis, you will need to know the number of observations, the mean, and the standard deviation from your sample data. We will first attempt to fit prior information on the standard deviation to an Inverse Chi distribution.

Considering your prior information, what do you feel is the most probable value of the standard deviation?

STDEV

4.9

Now, viewing your information about the standard deviation as if it were the result of an experiment of a certain sample size (N), how large an N-value do you feel your prior information to be worth? This N-value is the amount of weight that will be given to your prior knowledge relative to your present sample data,

N

6

The value of N and your estimate of the standard deviation imply that you would be willing to give even odds that the true standard deviation will lie in the 50% credibility interval from (3.9575) to (6.2858). This is the shortest possible 50% credibility interval for this value of N.

Inexperienced investigators are typically very conservative in their specification of N . When they see the length of the 50% credibility implied by their value, they typically increase this value.

Consider the above interval. It may or may not be satisfactory to you. You may increase the interval by selecting a smaller N, or decrease the interval by selecting a larger N. However, such changes will affect the relative weight you will be giving to your prior information. If you think you may be able to improve this interval by selecting a new value of N, please type that value. Otherwise please type a 0.

ans
20

The new 50% credibility interval, based on an N value of (20), extends from (4.3994) to (5.5032).
The old 50% credibility interval, based on an N value of (6), extends from (3.9575) to (6.2858).
If you think you may improve this interval by selecting a new value of N, please type that value. Otherwise please type a 0.

ans
25

The new 50% credibility interval, based on an N value of (25), extends from (4.4521) to (5.4284).
The old 50% credibility interval, based on an N value of (20), extends from (4.3994) to (5.5032).
If you think you may improve this interval by selecting a new value of N, please type that value. Otherwise please type a 0.

ans
0

On the basis of the information so far given, the prior marginal distribution on the standard deviation is an Inverse Chi with (24.0) degrees of freedom and a scale parameter of (24.5000).
This distribution has a mode of (4.9000), a median of (5.0720), and a mean of (5.1651).
A 50% credibility interval extends from (4.4521) to (5.4284).

> The value of N typically settled on is one that implies a somewhat longer interval than the investigator would have initially specified. By requiring the investigator to balance two judgments against each other, we hope to arrive at a coherent and realistic assessment.

If you think that you may be able to obtain a better prior distribution by respecifying your modal estimate of the standard deviation please type your new estimate, otherwise type 0

> This is one of several points at which the investigator can change his mind about some aspect of his prior beliefs. This is a particularly useful point. Most investigators are inappreciative of the asymmetry of the Inverse Chi density, and thus, are surprised by the length above the mode required to obtain a 50% HDR interval. A typical desire, at this point, is to lower the modal estimate slightly.

ans

4.7

Now, viewing your information about the standard deviation as if it were the result of an experiment of a certain sample size (N), how large an N-value do you feel your prior information to be worth? This N-value is the amount of weight that will be given to your prior knowledge relative to your present sample data,

N

20

The value of N and your estimate of the standard deviation imply that you would be willing to give even odds that the true standard deviation will lie in the 50% credibility interval from (4.2198) to (5.2786). This is the shortest possible 50% credibility interval for this value of N.

∞

Consider the above interval. It may or may not be satisfactory to you. You may increase the interval by selecting a smaller N, or decrease the interval by selecting a larger N. However, such changes will affect the relative weight you will be giving to your prior information. If you think you may be able to improve this interval by selecting a new value of N, please type that value. Otherwise please type a 0.

ans

0

On the basis of the information so far given, the prior marginal distribution on the standard deviation is an Inverse Chi with (19.0) degrees of freedom and a scale parameter of (21.0190). This distribution has a mode of (4.7000), a median of (4.9090), and a mean of (5.0245). A 50% credibility interval extends from (4.2198) to (5.2786).

If you think that you may be able to obtain a better prior distribution by respecifying your modal estimate of the standard deviation please type your new estimate, otherwise type 0

ans

0

Now considering your prior information about the mean and assuming that your modal estimate of the standard deviation is the actual true value, what is your best estimate of the true value of the unknown mean?

MEAN

21

Suppose I estimated a certain mean to be 100 and was willing to give even odds that the true mean would be between 95 and 105, then my central 50% credibility interval would be 10 units long. How long would a central interval around your estimate of the mean for the present problem need to be for you to be willing to give even odds that the true mean will lie in this interval?

INT

2

Again thinking of your prior knowledge as if it were the result of an experiment of a certain sample size (M), how many observations do you feel your prior knowledge on the mean to be worth?

M

5

The value of M you selected implies a 50% credibility interval which extends from (19.5823) to (22.4177).
The original interval you specified implies an M value of (10), and extends from (20.0000) to (22.0000).

Having treated the investigator gently in the first part of the program, we now attempt to show him that coherence in prior specifications is difficult to attain. The present investigator displays a rather high degree of coherence at this point compared with others who have used this program.

Please compare these two intervals. If your original interval is longer than the interval resulting from your selection of M, you may wish to select a smaller value of M. If your interval is smaller, you may wish to select a larger value of M. However, such changes will affect the relative weight you will be giving to your prior information. If you are satisfied with the interval (19.5823, 22.4177), M=(5), please type a 0, otherwise select a new value for M. Presumably you will want to select some value between M=(5) and M=(10), although you are free to select any value you wish.

ans
8

The new interval around your estimated mean is from (19.8792) to (22.1208), M=(8).
Your original interval around the mean was from (20.0000) to (22.0000), M=(10).
If you think you can improve your present interval by selecting a new value for M, please specify that value, otherwise please type a 0.

ans
10

The new interval around your estimated mean is from (19.9975) to (22.0025), M=(10).
Your original interval around the mean was from (20.0000) to (22.0000), M=(10).
If you think you can improve your present interval by selecting a new value for M, please specify that value, otherwise please type a 0.

ans
0

Your prior conditional distribution for the mean (given that the true standard deviation is 4.7000) is Normal, with a mean of (21.0000), a variance of (2.2090), and a standard deviation of (1.4863).
The 50% credibility interval around the mean extends from (19.9975) to (22.0025).

Your prior marginal distribution on the mean is a t distribution on (19) degrees of freedom with mean, mode and median (21.0000), and standard deviation (1.6121).
The 50% credibility interval about the mean extends from (19.9509) to (22.0491).

The investigator should note here that the t interval on the mean is longer than the normal interval.

Your final N value in fitting a prior on the standard deviation is (20.000)

Your final M value in fitting a prior distribution on the mean is (10.0000)

Your larger N value implies that you have more prior information about the standard deviation than about the mean. This does happen in some experiments. However, in many experimental situations the amount of prior information we have on the population parameters is the same. Hence, M and N would have the same value. If you wish, you may select a compromise value for N and M, between 6 and 51, which will be jointly consistent with your prior knowledge on both parameters. If the final result of this compromise is unsatisfactory you will still have the option of performing the posterior analysis with N and M unequal. If you do not wish to select a compromise value for N and M, please type a 0, otherwise type the compromise value.

ANS

15

Note that we offer the investigator the option of adopting a simplifying mode of analysis, but we do not force him into it. Our present experience suggests that in many situations, investigators feel that they have more information about the standard deviation than about the mean.

The compromise value of N and M selected results in the following distributions.

The prior marginal distribution on the population standard deviation is an Inverse Chi with (14.0) degrees of freedom, and a scale parameter of (18.2030). The mean of this distribution is (5.1486), the mode is (4.7000), and a 50% credibility interval extends from (4.1449) to (5.3910).

The prior conditional distribution on the mean is Normal with a mean of (21.0000), a variance of (1.4727), and a standard deviation of (1.214). The 50% credibility interval about the mean is from (20.1815) to (21.8185).

Your prior marginal distribution on the mean is a t distribution on (14.0) degrees of freedom,
with mean, median, and mode (21.0000), and standard deviation of (1.3568).
The 50% credibility interval about the mean extends from (20.1308) to (21.8692).

If you think you can obtain a more satisfactory compromise, please select a new value for NM.
Otherwise, please type a 0.

ans

16

The compromise value of N and M selected results in the following distributions.

The prior marginal distribution on the population standard deviation is an Inverse Chi with (15.0) degrees of freedom, and a scale parameter of (18.8000). The mean of this distribution is (5.1167), the mode is (4.7000), and a 50% credibility interval extends from (4.1628) to (5.3635).

The prior conditional distribution on the mean is Normal with a mean of (21.0000), a variance of (1.3806), and a standard deviation of (1.175). The 50% credibility interval about the mean is from (20.2075) to (21.7925).

12

Your prior marginal distribution on the mean is a t distribution on (15.0) degrees of freedom,
with mean, median, and mode (21.0000), and standard deviation of (1.3035).
The 50% credibility interval about the mean extends from (20.1614) to (21.8386).

If you think you can obtain a more satisfactory compromise, please select a new value for NM.
Otherwise, please type a 0.

ans

0

If you do not now wish to do a posterior analysis type 0
If you wish to do a posterior analysis using the compromise value for N and M type 1
If you wish to do a posterior analysis using the unequal values for N and M type 2

ans

2

If the investigator is still unhappy with his prior distribution and wants to abort the present analysis, he can depress the attention key. The program will note the point at which it was aborted and shift the program into the executive mode. The investigator can then start the program over at the beginning by typing "xeq" and depressing the return key. If he wishes to suppress printing of the introductory text, he can instead type "go to STDEV". If he is prepared to provide definite values for M and the standard deviation, he will instead type "go to MEAN".

POSTERIOR ANALYSIS WITH UNEQUAL N AND M

Having specified prior distributions on the standard deviation and on the mean, what is the number of observations of your present sample?

NUM

25

What is the mean of your sample?

SMEAN

21.3

What is the standard deviation from your sample?

SSD

5.2

The joint posterior mode for the mean and the standard deviation is (21.2143, 4.9309).
This is the point on the plane around which the probability is most highly concentrated.

The posterior marginal density on the population standard deviation is an inverse Chi distribution on (44.0) degrees of freedom with scale factor (33.4431). The posterior marginal mode is (4.9854), the posterior median is (5.0804), and the posterior mean is (5.1299)
The 50% credibility interval is from (4.6443) to (5.3703)

The posterior marginal distribution on the population mean is a t distribution on (44.0) degrees of freedom. It has a mean, median, and mode of (21.2143), and a standard deviation of (.8723). The 50% credibility interval extends from (20.6345) to (21.7940).
If you wish to reanalyze your data using the compromise value NM, please type a 1.
Otherwise type a 0.

ans

1

We note that the degrees of freedom for the t distribution depends on NUM and N (rather than M). The surprising sample value is SSD = 5.2 . While this value is inside the 50% HDR , it is .5 higher than the original modal estimate.

POSTERIOR ANALYSIS WITH N = M

Having specified prior distributions on the standard deviation and on the mean, what is the number of observations of your present sample?

NUM

25

What is the mean of your sample?

SMEAN

21.3

What is the standard deviation from your sample?

SSD

5.2

The joint posterior mode for the mean and the standard deviation is (21.1829, 4.9529).
This is the point on the plane around which the probability is most highly concentrated.

The posterior marginal density on the population standard deviation is an inverse Chi distribution on (40.0) degrees of freedom with scale factor (32.0986). The posterior marginal mode is (5.0130), the posterior median is (5.1181), and the posterior mean is (5.1732)
The 50% credibility interval is from (4.6540) to (5.4205)

The posterior marginal distribution on the population mean is a t distribution on (40.0) degrees of freedom. It has a mean, median, and mode of (21.1829), and a standard deviation of (.8132).
The 50% credibility interval extends from (20.6433) to (21.7226).

> There appears to be no important difference between using the values N = 20 and M = 10 and the compromise value NM = 16 . Should this result be observed consistently, we would feel confident in reserving our use of the more complicated unequal NM analysis to critical situations.

If you do not wish to analyze further data sets, please type a 0.
If you wish to analyze further data sets using your present prior distributions, please type a 1.
If you wish to respecify your prior distributions, please type a 2.

Ans

2

> At this point, a second investigator, having somewhat different prior beliefs, is invited to quantify his beliefs so that a new posterior analysis can be done, and the result compared with that of the original analysis. The new investigator is more knowledgeable about educational testing, the state of Iowa and ACT scores, and more experienced with CADA.

Considering your prior information, what do you feel is the most probable value of the standard deviation?

STDEV

5.3

Now, viewing your information about the standard deviation as if it were the result of an experiment of a certain sample size (N), how large an N-value do you feel your prior information to be worth? This N-value is the amount of weight that will be given to your prior knowledge relative to your present sample data, and should be some integer value greater than or equal to 6, but less than or equal to 51. These limits will be satisfactory in most applications.

N

20

The value of N and your estimate of the standard deviation imply that you would be willing to give even odds that the true standard deviation will lie in the 50% credibility interval from (4.7585) to (5.9525). This is the shortest possible 50% credibility interval for this value of N.

Consider the above interval. It may or may not be satisfactory to you. You may increase the interval by selecting a smaller N, or decrease the interval by selecting a larger N. However, such changes will affect the relative weight you will be giving to your prior information. If you think you may be able to improve this interval by selecting a new value of N, please type that value. Otherwise please type a 0.

ans

0

On the basis of the information so far given, the prior marginal distribution on the standard deviation is an Inverse Chi with (19.0) degrees of freedom and a scale parameter of (23.7023). This distribution has a mode of (5.3000), a median of (5.5357), and a mean of (5.6659). A 50% credibility interval extends from (4.7585) to (5.9525).

If you think that you may be able to obtain a better prior distribution by respecifying your modal estimate of the standard deviation please type your new estimate, otherwise type 0

ans

0

Now considering your prior information about the mean and assuming that your modal estimate of the standard deviation is the actual true value, what is your best estimate of the true value of the unknown mean?

MEAN

21.5

Suppose I estimated a certain mean to be 100 and was willing to give even odds that the true mean would be between 95 and 105, then my central 50% credibility interval would be 10 units long. How long would a central interval around your estimate of the mean for the present problem need to be for you to be willing to give even odds that the true mean will lie in this interval?

INT

1

Again thinking of your prior knowledge as if it were the result of an experiment of a certain sample size (M), how many observations do you feel your prior knowledge on the mean to be worth?

M

20

The value of M you selected implies a 50% credibility interval which extends from (20.7006) to (22.2994).
The original interval you specified implies an M value of (51), and extends from (21.0000) to (22.0000).

Please compare these two intervals. If your original interval is longer than the interval resulting from your selection of M, you may wish to select a smaller value of M. If your interval is smaller, you may wish to select a larger value of M. However, such changes will affect the relative weight you will be giving to your prior information. If you are satisfied with the interval (20.7006, 22.2994), M=(20), please type a 0, otherwise select a new value for M. Presumably you will want to select some value between M=(20) and M=(51), although you are free to select any value you wish.

ans

0

Your prior conditional distribution for the mean (given that the true standard deviation is 5.3000)
is Normal, with a mean of (21.5000), a variance of (1.4045), and a standard deviation of (1.1851).
The 50% credibility interval around the mean extends from (20.7006) to (22.2994).

Your prior marginal distribution on the mean is a t distribution on (19) degrees of freedom
with mean, mode and median (21.5000), and standard deviation (1.2854).
The 50% credibility interval about the mean extends from (20.6635) to (22.3365).
If you do not now wish to do a posterior analysis type 0
If you wish to do a posterior analysis type 1

ans

1

POSTERIOR ANALYSIS WITH N = M

Having specified prior distributions on the standard deviation and on the mean, what is the number of observations of your present sample?

NUM

25

What is the mean of your sample?

SMEAN

21.3

What is the standard deviation from your sample?

SSD

5.2

The joint posterior mode for the mean and the standard deviation is (21.3889, 5.1883).
This is the point on the plane around which the probability is most highly concentrated.

The posterior marginal density on the population standard deviation is an inverse Chi distribution on (44.0) degrees of freedom with scale factor (35.1887). The posterior marginal mode is (5.2456), the posterior median is (5.3455), and the posterior mean is (5.3977)
The 50% credibility interval is from (4.8868) to (5.6506)

The posterior marginal distribution on the population mean is a t distribution on (44.0) degrees of freedom. It has a mean, median, and mode of (21.3889), and a standard deviation of (.8094).
The 50% credibility interval extends from (20.8509) to (21.9269).

> The second investigator had somewhat higher prior estimates both for the mean and for the standard deviation. However, the posterior distributions for the two investigators are very much closer together than were the prior distributions. The different prior distributions do have an effect on the two posterior distributions, but it is not large--considering the small size of the available sample.

19

If you do not wish to analyze further data sets, please type a 0.
If you wish to analyze further data sets using your present prior distributions, please type a 1.
If you wish to respecify your prior distributions, please type a 2.

Ans
1

If you do not now wish to do a posterior analysis type 0
If you wish to do a posterior analysis type 1

ans

> The second investigator does not feel that his posterior distributions are sufficiently tight. He has the option of obtaining another sample of 25 and wonders whether that is likely to be sufficient for his purposes. He, therefore, asks for the computer to determine what his posterior distribution would be if his second sample were such that the total sample had the same mean and standard deviation as the first sample.

POSTERIOR ANALYSIS WITH N = M

Having specified prior distributions on the standard deviation and on the mean, what is the number of observations of your present sample?

NUM

50

What is the mean of your sample?

SMEAN

21.3

What is the standard deviation from your sample?

SSD
5.2

The joint posterior mode for the mean and the standard deviation is (21.3571, 5.1926).
This is the point on the plane around which the probability is most highly concentrated.

The posterior marginal density on the population standard deviation is an inverse Chi distribution on (69.0) degrees of freedom with scale factor (43.7535). The posterior marginal mode is (5.2295), the posterior median is (5.2929), and the posterior mean is (5.3255)

The posterior marginal distribution on the population mean is a t distribution on (69.0) degrees of freedom. It has a mean, median, and mode of (21.3571), and a standard deviation of (.6389). The 50% credibility interval extends from (20.9302) to (21.7841).

If you do not wish to analyze further data sets, please type a 0.

If you wish to analyze further data sets using your present prior distributions, please type a 1.

If you wish to respecify your prior distributions, please type a 2.

Ans

The second investigator now wants to know what the first investigators posterior distribution will be--given these same 50 observations. He, therefore, shifts the program to executive mode and types "go to post". The computer then requests values for NM, MEAN, and STDEV. Given these values, the computer computes the posterior distribution using the compromise value of NM.

21

690. WAS LAST LINE EXECUTED

go to post

NM

16

MEAN

21

STDEV

4.7

POSTERIOR ANALYSIS WITH N = M

Having specified prior distributions on the standard deviation and on the mean, what is the number of observations of your present sample?

NUM

50

What is the mean of your sample?

SMEAN

21.3

What is the standard deviation from your sample?

SSD

5.2

The joint posterior mode for the mean and the standard deviation is (21.2273, 5.0468).
This is the point on the plane around which the probability is most highly concentrated.

The posterior marginal density on the population standard deviation is an inverse Chi distribution on (65.0) degrees of freedom with scale factor (41.3102). The posterior marginal mode is (5.0849), the posterior median is (5.1504), and the posterior mean is (5.1841)

The posterior marginal distribution on the population mean is a t distribution on (65.0) degrees of freedom. It has a mean, median, and mode of (21.2273), and a standard deviation of (.6406). The 50% credibility interval extends from (20.7994) to (21.6552).

22

Both investigators now feel that some off-line discussions may be profitable. Should further exploration be required, it will be easy to do this with the CADA program.

If you do not wish to analyze further data sets, please type a 0.
If you wish to analyze further data sets using your present prior distributions, please type a 1.
If you wish to respecify your prior distributions, please type a 2.

Ans

0

This completes the program of Bayesian analysis for the Normal model with unknown mean and standard deviation.

The Interactive System for Bayesian Analysis of Educational Data that you have been using was developed under the sponsorship of the Iowa Testing Program(ITP), and the American College Testing Program(ACT), with the support of the University of Iowa Computer Center(UCC).

PROJECT STAFF

Sequence Programming

VICTOR J. ORMSBY

Systems Programming

DAVID E. CHRIST

Mathematical Development

PAUL H. JACKSON

Coordinator for U.C.C.

JAMES WHITELY

Principal Investigator and Project Director

MELVIN R. NOVICK

** 740. EXECUTION STOP.

-

The exploration described in this example can now be done at the University of Iowa Computer Center in less than 30 minutes and at a cost of less than four dollars. Further economies will be forthcoming when the program scripts are transferred from slow core to disk storage. Even so, with the present technology, costs will remain sufficiently high to preclude use of the system on a routine basis.

The present technology is somewhat unsatisfactory in another way. Sitting at the terminal, an investigator with even very limited exposure to the system feels that he is spending much of his time waiting for the typewriter to print out text that he saw printed several times earlier in the analysis. He needs to have this text to refer to, but he finds that a large percentage of his time is spent watching the typewriter. If this problem could be solved, the exploration described above could probably be done in 15 rather than 30 minutes. The resulting decrease in connect time charges and other economies could be expected to reduce the cost of above analysis to about two dollars.

The discussion in the previous paragraph is only one part of the larger problem concerned with the interface of the investigator with the computer. The standard problem has always been that of making it easy enough for investigators to use the computer so that they felt this use was, in fact, saving them time and effort. Changing "input-output" procedures, inadequate program write-ups, and poorly labeled output formats are continuing computer center problems that can be solved and often are solved--but not always. Given adequate interest and time, computer programmers can produce output formats that are easy to read and; indeed, using the kind of annotation exhibited in the paper, they

can, and have, produced useful sample problems. Unfortunately, this kind of work is seldom given high priority and, hence, is seldom done.

Another continuing problem of statistical computing is that of turn-around time which can range from 15 minutes to a day or more depending on the demands being placed on a particular computer. At a computer center that is efficiently run and not overtaxed, an investigator may be able to get five or six runs per day either by physically carrying his input to the computer center or submitting his jobs through remote batch-job-entry. This, however, is inefficient. Each batch submission takes time to prepare and deliver, and this costs money. (Our accounting practices do not place any charge on student time misspent, but the costs are real.) In a work environment, accounting practices are more realistic. In the example we have presented, a minimum of six runs would be required to get out the full Bayesian analysis. Assuming no mishaps (an event having probability much less than one under ideal conditions), one could hardly expect to get this work done in less than a day. Indeed, one would not do this by computer. It would be too costly.

A skilled worker could do the job in a couple of hours, but that extra hour and forty-five minutes of work is a poor alternative to the two dollars of projected computer cost. More importantly, what of the investigator who has not maintained current computer skill? Such a person would probably spend a day on the project, with or without the computer and then, with high probability, not get it right.

For the reasons given above and for others not touched upon here, it is certain that interactive computing will be the standard computing mode in the not too distant future. Given the acceptance of this conclusion, it then becomes

clear that a beginning should be made in the study of the interface between man and computer in an interactive environment.

An admitted difficulty of this study is that hardware technology is changing so swiftly that by the time any evaluation and development is complete, it will be less than completely up to date. Nevertheless, certain projections concerning hardware development and costs can be made with some confidence, and one of these will be touched on shortly.

The basic problem of giving the investigator the information and guidance he needs without continually reprinting a book is easily solved by the use of a cathode ray tube terminal. With at least one currently available terminal, printing can be done at 240 characters per second, as compared to 15 characters per second for a typewriter terminal. Second, basic instructions and guidelines can remain on the screen (in half tone if desired) with only investigator responses and computations therefrom changing. Thus, investigators can read or ignore background information according to their needs.

At present, cathode ray tube terminals cost about $3,000 each. Within three years improved versions of the same terminals will be available in lots of three or more at half that price. Either of these figures is less than the cost of purchasing additional typewriter terminals. There are other equipment and personnel costs involved in interfacing the terminals to a large computer. These are not small but are largely of a one-time nature. One approach involves the use of a mini-computer to accomplish the interface.

Teaching and use of statistics in education has changed little in the last 15 years. The computer has made computational work easier and has made

it possible for more and more sophisticated work to be done. Computer training, however, has always been separate from statistics training under the assumption that when the Ph.D. candidate got to his dissertation, he could learn what he needed to know about programming to get his thesis done.

What is being proposed here is that interaction with the computer be phased into the teaching of statistics at all but, perhaps, the most elementary level. Thus, almost from the beginning, the student will learn to use the computer in an efficient and effective manner, and he will be learning to use the computer in precisely the way that he will use it on the job.

The interactive program exhibited here was developed over the last year, largely through a trial and error process. Refinements of the text and changes in the programming are introduced whenever a meaningful improvement seems possible. The example presented in this paper was prepared for presentation at a seminar sponsored by The International Statistical Institute and held at The Health Sciences Computing Facility in Los Angeles, August 27-September 1, 1971. I am indebted to the organizers and participants of this conference both for specific comments and for general encouragement.